REMARKS

Applicant respectfully requests reconsideration and allowance of the subject application in view of the amendments and the remarks to follow. Claims 1-48 are pending in this application.

Request for Return of Form PTO-1449

Applicant notes that copies of the PTO-1449 forms tendered on April 10, 2004, have been properly initialed and returned to Applicant. Applicant also notes that one of the two PTO-1449 forms tendered on August 2, 2004 was also properly initialed and a copy thereof returned to Applicant.

Accordingly, Applicant is resending the PTO-1449 listing two-nonpatent references herewith, and respectfully requests that the references listed on both that PTO-1449 and the PTO-1449 form tendered on October 8, 2004 be considered, and that properly initialed and dated copies thereof be returned to Applicant.

Summary of Unpatentability Rejections Under 35 U.S.C. § 103(a)

Rejection I.

Claims 1, 2, 4-7 and 9-18 are stated (p. 2, item 3) to stand rejected under 35 U.S.C. §103(a) as being unpatentable over Duvillier et al., published U.S. Patent Application No. 2002/0074082 (hereinafter "Duvillier") in view of Zakai et al., U.S. Patent No. 6,415,372 (hereinafter "Zakai").

Rejection II.

Claims 3 and 8 are stated (p. 7, item 4) to stand rejected under 35 U.S.C. §103(a) as being unpatentable over Duvillier in view of Zakai and further in view of Mukherjee et al., U.S. Patent No. 6,466,978 (hereinafter "Mukherjee").

Rejection III.

Claims 19, 22-33, 35-37 and 39-41 are stated (p. 8, item 5) to stand rejected under 35 U.S.C. §103(a) as being unpatentable over Zakai.

Rejection IV.

Claims 20 and 38 are stated (p. 16, item 6) to stand rejected under 35 U.S.C. §103(a) as being unpatentable over Zakai in view of Rabinovich, U.S. Patent No. 6,484,204 (hereinafter "Rabinovich").

Rejection V.

Claims 21 and 34 are stated (p. 17, item 7) to stand rejected under 35 U.S.C. §103(a) as being unpatentable over Zakai in view of Mukherjee.

Rejection VI.

Claims 42-44 and 46-48 are stated (p. 18, item 8) to stand rejected under 35 U.S.C. §103(a) as being unpatentable over Falls (no patent number provided in Office Action, Applicant assumes that what was intended was U.S. Patent No. 5,991,771; clarification is respectfully requested) in view of Zakai.

Rejection VII. Claim 45 is stated (p. 21, item 9) to stand rejected under 35 U.S.C. §103(a) as being unpatentable over Falls in view of Zakai and further in view of Rabinovich. Applicant respectfully disagrees and requests reconsideration.

Traverse I.

Claims 1, 2, 4-7 and 9-18 stand rejected as being unpatentable over Duvillier in view of Zakai. Duvillier is directed (Title) to a: "System modification processing technique implemented on an information storage and retrieval system". Duvillier teaches that: "A technique is disclosed for implementing system modification operations in an information storage and retrieval system. The information storage and retrieval system includes persistent memory configured or designed to store object data. The persistent memory includes at least one data file for storing object data. A first system modification request relating to a first data file is received, the first data file including a first object stored therein. The first system modification request is then implemented. According to a specific embodiment, the implementation of the first system modification request includes suspending write access to the first data file. Concurrently, while the first system modification request is being implemented, updated information relating to the first object may be stored in the persistent memory. According to a specific embodiment, the information storage and retrieval system corresponds to a non-positional, non-log based information storage and retrieval system. According to different embodiments, the information storage and retrieval system of the present invention may be configured to handle a variety of different system modification requests, including, for example, a request to add a mirror data file to be associated with a primary data file, a request to take the primary data file off-line, a request to take the mirror data file off-line. Moreover, according to a specific implementation, the implementing of the first

system modification request may performed in real-time, without blocking access to object data stored in the persistent memory." (Abstract).

In contrast, claim 1 recites "A method for determining where to store object replicas, the method comprising: receiving an indication of a homeless replica of an object, wherein the object has a plurality of replicas including the homeless replica; determining an initial placement for the homeless replica, wherein the initial placement is one of a plurality of devices in a system; evaluating whether any replicas of the object can be swapped with one of a plurality of replicas of another object and not reduce a combined object availability of the two objects; and swapping a replica of the object with the one of the plurality of replicas of the other object only if the swapping does not reduce the combined object availability of the two objects", which is not taught, disclosed, suggested or motivated by the cited references, alone or in any proper combination.

The Office Action cites (page 2, item 3) page 8, right col., lines 36-53 of Duvillier relative to "receiving an indication of a homeless replica ..." in claim 1. Applicant infers that the Office Action intends to refer to the description of Fig. 6 contained on that page, i.e., ¶ 93 et seq. That passage and Fig. are concerned with swapping of a mirror file for a primary file.

In contrast, Applicant describes a "homeless" object. For example, the specification, at page 4, lines 13-18, Applicant states that: "Determining the location where the file replicas should be stored (that is, on which devices 102 the file replicas should be stored) is a two part process. The first part, referred to as initial placement, determines on which device 102 a particular file replica should

be stored when the replica is not currently stored anywhere or its previous storage location is no longer available (in other words, the replica is "homeless")."

Further clarification of the meaning of the term "homeless" in the context of Applicant's claims is provided at page 10, line 14 et seq., with reference to Fig. 3, stating that: "Fig. 3 is a flowchart illustrating an exemplary process 200 for placing file replicas in a system. Process 200 is implemented by one or more devices responsible for maintaining one or more directories. Process 200 is idle (act 202) until one of two events occurs – a homeless replica needs placement, or a replica swap is triggered. When a homeless replica exists, the homeless replica needs placement and an initial placement part of process 200 is invoked. A replica may be homeless for a variety of different reasons, such as creation of a new file, a device previously storing a replica is removed from the system (e.g., to no longer be part of distributed file system 160 of Fig. 2, or because it has broken), a replica was evicted from a device because space was needed on the device for local storage, the number of replicas for the file has increased, and so forth."

This clearly is not equivalent to swapping of a mirror file for a primary file, as described by Duvillier, each of which has a well-defined storage location and thus is not homeless. More specifically, Duvillier teaches (page 1, right hand col., 12 et seq.) that:

> For example, one limitation of conventional RDBMS systems is that the relational nature of objects stored in the RDBMS system requires that all updates to a data object stored within the database 120 be written each time to **the same physical location** (e.g. disk space) where that object is stored. Because of this requirement, such systems are typically referred to as "positional" database systems since it is important that the relative position of each object stored in the database be maintained in order for the relational database to function properly. Moreover, when updates are being performed on portions of data stored within a positional database system, users

> will typically be unable to access any portion of the updated data until after the entirety of the data update has been completed. If the user attempts to access a portion of the data while the update is occurring, the user will typically experience a hanging problem, or will be handed dirty data (e.g. stale data) until the update transaction(s) have been completed.
>
> In light of this problem, content providers typically resort to setting up a second data file (typically referred to as a "mirror" data file) which includes an identical copy of the information stored in the original or "primary" data file. In this way users are provided with the ability to access desired information from the mirror data file at times when the primary data file is off-line. However, it will be appreciated that such an approach demands a relatively large amount of resources for implementation, particularly with respect to memory resources.

In other words, Duvillier is concerned with an environment, such as a banking computer system, in which a high degree of reliability is expected and in fact demanded (see, e.g., ¶ 8 on page 1). For this reason, Duvillier operates in the context of a system that employs "atomic writes" to memory, that is, either the write is completed and accurate, or the write does take place - thus avoiding a "torn write" with resultant corrupted data.

As a result, when one aspect of the system is accessing a portion of memory, other portions of the system are excluded - hence the need for a mirror memory location and a system for tracking which of the two is most current and which is "stale". Duvillier not only teaches a system in which there are no "homeless" objects, but a system where objects have fixed physical addresses and **never** move.

Duvillier teaches a system in which both the primary and mirror locations may be "clean", i.e., represent current data, both may be "dirty" and thus need to be written (¶ 85 describes a "synchronous write" to update both at substantially the same time) or a situation where it is known that one or the other represents "stale"

or "dirty" data. Further, the mirror memory location may be needed when the primary data location corresponds to a disc that is being physically removed from the system (see ¶ 14)

Swapping between "stale" data represented at a primary location and a current data represented at a mirror location is completely non-sequitur to the subject matter of any of Applicant's claims and hardly results in either such location becoming suddenly "homeless". In fact, Duvillier is void of the term "homeless".

Zakai fails to cure the deficiencies of Duvillier. Zakai is directed (Title) to: "Rolling back storage subsystem reconfigurations". Zakai describes: "A method and an apparatus for reconfiguring a storage subsystem by performing an ordered sequence of reconfigurations of physical storage volumes of the storage subsystem. The method and apparatus perform a portion of the sequence of reconfigurations, in response to receiving a rollback request, in an order that is reversed with respect to the order of the sequence." (Abstract).

More specifically, neither Duvillier nor Zakai are concerned with anything to do with homeless objects. Each is completely void of the term "homeless". As such, it is inconceivable that combining their disclosures could provide the subject matter of any of Applicant's claims.

Zakai describes swapping of entire physical storage volumes (e.g., A, B etc., Fig. 1A) in order to balance accessions of distinct memory banks (e.g., 15, 16, 17). This is clearly set forth at col. 5, line 60 et seq., stating that:

> For each identified pair of storage devices 15-17, the service processor 28 searches for swaps of physical storage volumes A-G that produce better balanced workloads (step 59). A swap must reduce historical imbalances of workloads for the paired storage

devices 15-17 by more than predetermined threshold amount, e.g., ten percent, to be retained. **Two storage volumes A-G must have the same size and emulation characteristics to qualify as potentially swappable**.

Attempting to combine the teachings of Duvillier with those of Zakai clearly renders each of them unsuitable for their intended purpose. Duvillier teaches a system where files have permanent physical addresses. Zakai teaches swapping of physical storage volumes and thus of explicit modification of physical addresses. It is inappropriate to employ references in a manner that renders their teachings unsuitable for their intended purposes. This is explained below in more detail with reference to MPEP §2143.01.

In a subsection entitled "THE PROPOSED MODIFICATION CANNOT RENDER THE PRIOR ART UNSATISFACTORY FOR ITS INTENDED PURPOSE", this MPEP portion states that: "If proposed modification would render the prior art invention being modified unsatisfactory for its intended purpose, then **there is no suggestion or motivation to make the proposed modification**. *In re Gordon*, 733 F.2d 900, 221 USPQ 1125 (Fed. Cir. 1984)."

Accordingly, the rejection of claims 1, 2, 4-7 and 9-18 is prima facie defective and should be withdrawn, and claims 1, 2, 4-7 and 9-18 should be allowed.

Traverse II.

Claims 3 and 8 stand rejected as being unpatentable over Duvillier in view of Zakai and further in view of Mukherjee. Mukherjee is directed (Title) to: "Multimedia file systems using file managers located on clients for managing network attached storage devices". Mukherjee describes: "A multi-media file system for conmmunicating [sic] information between a multi-media client and a network storage device over a network. The file system includes a cluster that comprises one cluster manager and at least one file manager with each network storage device. The cluster manager is located on a client, includes an admission controller for controlling the admission of a request from a client for a file operation upon a selected file. A network bandwidth request from the admission controller is responded to by a network status determiner included in the cluster manager. The network status determiner determines the available network bandwidth. Each file manager is located on one of the clients. The file managers manage file maintenance procedures of corresponding files located on the network storage device. Each file manager includes a disk status determiner for determining the available disk bandwidth. The disk status determiner responds to a request from the admission controller." (Abstract).

In contrast, claim 3 recites "A method as recited in claim 1, wherein the method is implemented by multiple computing devices in a serverless distributed file system", while claim 8 recites "A method as recited in claim 1, wherein the object comprises a directory", which recitations are not taught, disclosed, suggested or motivated by Duvillier, with or without Zakai and/or Mukherjee.

As noted above, Duvillier and Zakai are non sequitur to the subject matter recited in claim 1 and thus claims dependent therefrom. Mukherjee fails to cure the deficiencies of Duvillier and Zakai.

Mukherjee is specifically concerned with problems unique to sharing of multimedia data files including both audio and video information (see, e.g., Title, Abstract etc.). Mukherjee points out that multimedia data is sensible only when played back in temporal contiguity (col. 1, line 26 et seq.) and that this places increased constraints on storage and retrieval of such information, including increased bandwidth requirements. As such, Mukherjee is concerned with completely different kinds of data handling problems than those being addressed by Duvillier (e.g., textual banking records) or Zakai (relating to reconfiguration of a distributed storage system for unspecified data types, see, e.g., col. 1, line 53 et seq.).

As such, there can be no motivation to combine their teachings, and the result of such combination fails to provide the elements of Applicant's claims. For at least these reasons, the rejection of claims 3 and 8 is improper and should be withdrawn, and claims 3 and 8 should be allowed.

Traverse III.

Claims 19, 22-33, 35-37 and 39-41 stand rejected as being unpatentable over Zakai. In contrast to Zakai, claim 19 recites "One or more computer readable media having stored thereon a plurality of instructions that, when executed by one or more processors of a computing device, causes the one or more processors to perform acts comprising: working, in conjunction with one or more other processors of another computing device, to determine whether a replica of a file managed by the computing device and a replica of another file managed by the other computing device can be swapped with one another to bring an availability of the file and an availability of the other file closer; and swapping the replica of the file and the replica of the other file only if the swapping brings the availability of the file and the availability of the other file closer", which is not taught, disclosed, suggested or motivated by Zakai.

Also in contrast to Zakai, claim 33 recites "A serverless distributed file system comprising: a first plurality of computing devices storing files; a second plurality of computing devices managing storage of the files; wherein a first computing device of the second plurality of computing devices selects a file for which it manages storage and communicates with a second computing device of the second plurality of computing devices to determine whether a replica of the file and a replica of another file for which the second computing device manages storage can be swapped in order to improve a combined file availability; and if the replicas can be swapped to improve the combined file availability, then instructing the one of the first plurality of computing devices on which the replica of the file is stored to transfer the replica of the file to the one of the first plurality of

computing devices on which the replica of the other file is stored, and instructing the one of the first plurality of computing devices on which the replica of the other file is stored to transfer the replica of the other file to the one of the first plurality of computing devices on which the replica of the file is stored" and claim 36 recites "One or more computer readable media having stored thereon a plurality of instructions that, when executed by one or more processors of a computing device, causes the one or more processors to: initially place replicas of a file on different ones of a plurality of devices using a first process; and subsequently improve the placement of replicas of a plurality of files by evaluating whether any replicas of a first file can be swapped with any replicas of a second file without a reduction in the combined file availability of the first and second files, and swapping a replica of the first file with a replica of the second file if the swapping results in no reduction in the combined file availability of the first and second files", which is not taught, disclosed, suggested or motivated by the cited references.

The Office Action states (page 8, item 5) that Zakai teaches "working, in conjunction with one or more other processors of another computing device, to determine whether a replica of a file managed by the computing device and a replica of another file managed by the other computing device can be swapped with one another to bring an availability of the file [of the file and an availability of the other file closer]" as [sic] (col. 6, lines 5-15)" and "and swapping the replica of the file and the replica of the other file only if the swapping brings the availability of the file [and the availability of the other file closer]" as [sic] (col. 6, lines 5-15)." (portions elided in the Office Action are shown in square brackets) as recited in claim 19. Applicant disagrees, for several reasons.

First, the portions cited from claim 19 are disjoint to the point of incomprehensibility. Shattering a claim into unintelligible shards cannot possibly provide the claimed subject matter.

Second, the cited portion of Zakai reads as follows: "Data availability will not be compromised by the swaps. The service processor 28 makes a check to determine whether performing the swap will impact data availability. Availability is less affected if the storage subsystem 10 has multiple copies of the data to be swapped, e.g., in a redundant array of inexpensive disks (RAID). If multiple copies exist, the swap of a storage volume A-G holding one copy does not reduce the overall availability of the data. If swapping compromises data availability, the swap is not performed at step 60."

However, Zakai provides no measure or definition of the phrase "data availability", beyond noting that increased redundancy results in less of an effect on availability. Zakai provides no hint whatsoever that closeness of one file or another to anything is a parameter of concern. It is unclear whether Zakai is considering the effect of data being unavailable for a duration of a swap, because the data are being read from one physical volume and written to another, in a different storage device, or more generally, in terms of bottlenecking of a heavily used storage device that includes a data file.

In fact, Zakai is void of the term "close" or the term "closeness". As such, it is inconceivable that Zakai could suggest or motivate the subject matter of claim 19.

The Office Action further states (page 9) that "If multiple copies exist, the swap of a storage volume A-G holding one copy does not reduce the overall

availability of the data to Duvillier's system in order to reduce the network traffic that has to be generated to make a distribution decision, and reduces the complexity of such decision making." This passage has no discernible relationship at all to the rejection, to claim 19, to the subject matter of Zakai or even to the subject matter of Duvillier. The rejection of claim 19 does not involve Duvillier. Duvillier is not concerned with network traffic, and neither is Zakai, nor is claim 19. Clarification of the rejection is respectfully requested.

The Office Action states (page 12) that "As to claim 33, Zakai teaches the claimed limitations: "A first plurality of computing devices storing files" as host devices storing file [sic] (fig. 1A; col. 3, lines 3-23); "a second plurality of computing devices" as storage devices storing file [sic] (fig. 1A, col. 3, lines 3-15)". Applicant disagrees.

Col. 3, line 3 et seq. of Zakai states that: "The host devices 12-14 and storage devices 15-17 connect to the global memory 19 through channels 20' and 20, respectively, e.g., busses or networks." At col. 2, line 66 et seq., Zakai states that "FIG. 1A shows a storage subsystem 10 that provides shared data storage to a group of host devices 12-14. The shared data storage is located on physical storage devices 15-17, e.g., multiple disk devices, that are accessible to the host devices 12-14 through a global memory 19." Zakai clearly teaches that storage devices 15-17 comprise memories such as hard drives, and to not comprise computing devices. Computer parts are not equivalent to computers.

The Office Action states (page 14) that, with respect to claim 36: "Zakai teaches the claimed limitations: "initially place replicas of a file on different ones of a plurality of devices using a first process as [sic] (col. 6, lines 45-50); "and

subsequently improve the placement of replicas of a plurality of files by evaluating whether any replicas of a first file can be swapped with any replicas of a second file without a reduction in the combined availability of the first and second files" as [sic] (col. 6, lines 15-25)." Applicant disagrees and respectfully requests reconsideration.

As noted above, Zakai is unclear as to how the term "availability" is used in therein. If Zakai contemplates having three or more replicas of a file, availability is not as badly compromised during a file swap because at least one copy of the file should be available even if two others are "locked" as part of being swapped.

For at least these reasons, the rejection of claims 19, 22-33, 35-37 and 39-41 is inapposite and should be withdrawn, and claims 19, 22-33, 35-37 and 39-41 should be allowed.

Traverse IV.

Claims 20 and 38 stand rejected as being unpatentable over Zakai in view of Rabinovich. Rabinovich is directed (Title) to a: "System and method for allocating requests for objects and managing replicas of objects on a network". Rabinovich describes: "A system and method for distributing requests for objects to hosts that store replicas of the objects, and for managing the placement of the replicas among hosts. Metrics for the historical demand of a replica at a host and the distance of the host from the requester of the object are evaluated and used to make decisions as to where to forward the request substantially independently from any input provided by a host to which a request is forwarded. This simplifies autonomous replica placement decisions made by hosts. A host substantially autonomously uses request metric and load information to select a replica to be deleted, migrated or replicated, and to delete, migrate or replicate a selected replica." (Abstract).

In contrast to Rabinovich and/or Zakai, claim 20 recites "One or more computer readable media as recited in claim 19, wherein the swapping comprises communicating with a first device on which the replica of the file is stored and a second device on which the replica of the other file is stored to have the first device transfer the replica of the file to the second device and delete the replica of the file on the first device, and to have the second device transfer the replica of the other file to the first device and delete the replica of the other file on the second device", while claim 38 recites "One or more computer readable media as recited in claim 36, wherein swapping the replica of the first file with the replica of the second file comprises communicating with a first device on which the replica of

the first file is stored and a second device on which the replica of the second file is stored to have the first device transfer the replica of the first file to the second device and delete the replica of the first file on the first device, and to have the second device transfer the replica of the second file to the first device and delete the replica of the second file on the second device", which recitations are not taught, disclosed, suggested or motivated by Rabinovich in any proper combination with Zakai.

The Office Action states (page 16, item 6) that: "As to claim 20, Zakai teaches the claimed limitation "communicating with a first device on which the replica of the file is stored and a second device on which the replica of the other file is stored to have the first device transfer the replica of the file to the second device and delete the replica of the file on the first device, and to have the second device transfer the replica of the other file to the first device" as (col. 6, lines 45-60". Applicant disagrees.

This passage in Zakai is a portion of a description of Fig. 4A, beginning at col. 6, line 26, which passage states that:

> Referring again to FIG. 4A, the data of the table 46 will be used to illustrate one method for selecting data swaps in steps 57-59 of FIG. 5B.
>
> At step 57, the service processor 28 uses a method that ranks the storage devices 15-17 based on average workload per non-excluded time slice. Non-excluded time slices correspond to rows 1-3 of exemplary table 46 as is seen from the column 49. During the non-excluded time slices, table 46 shows that the storage devices 15 and 16 had 66 and 36 accesses, respectively. During the same period, the storage device 17 had 54 accesses (not shown in FIG. 4A). From the numbers of accesses, the average workloads of the storage devices 15, 16, and 17 are 22, 12, and 18 accesses per time slice, respectively. Thus, the service processor 28 will rank the respective storage devices 15, 16, and 17 as most busy, least busy, and second most busy, respectively.

> At step 58, the service processor 28 pairs off the most and least busy of the storage devices 15-17. Thus, the service processor 28 pairs storage devices 15 and 16 and determines that the storage device 17 will not participate in data swaps.
>
> At step 59, the service processor 28 uses a search methodology to select data swaps that decrease workload imbalances by more than a threshold amount. Since the imbalance between the storage devices 15 and 16 is ten accesses per time slice, only data swaps that reduce the imbalance by at least 2.5 accesses per time slice can be selected if the threshold is 25%. Swapping the data of storage volume A with storage volume D or E will reduce the imbalance between storage devices 15 and 16 by 1.33 and 2 accesses per time slice, respectively. Both reductions are too small for the service processor 28 to select these data swaps. Swapping storage volumes A and E will reduce the imbalance between the storage devices 15 and 16 by 4.66 accesses per time slice, which is greater than the exemplary threshold of twenty-five percent. Thus, the service processor 28 will select the data swap between volume A and volume E at step 59.
>
> After selecting the data swap between storage volumes A and E, the service processor 28 searches for other data swaps between the remaining storage volumes B, C and D, F. Any further selections of data swaps must further decrease the workload imbalance between the storage devices 15 and 16 by an above threshold amount. Swapping the storage volumes A and E makes the workload of the storage device 15 less than the workload of the storage device 16. Any other data swaps between the storage volumes B, C and the storage volumes D, F will increase the imbalance between the storage devices 15 and 16. Thus, the method 56 will only select to swap the storage volumes A and E for the workloads shown in FIG. 4A.

In other words, Zakai describes a very specific technique for swapping entire contents of physical storage volumes among a number of storage devices, responsive to very carefully-defined demand-based rules. Zakai also describes a mechanism whereby such is coordinated by a central device, the service processor 28.

In strong contrast, Rabinovich is concerned with a totally different set of problems within the narrowly-defined scope of problems relating to replication of

files, or "mirroring", in the context of Internet scalability (col. 1, line 15 et seq.). Rabinovich teaches (col. 1, line 31 et seq.) that "It is important to properly decide where to store the replicas, and how to allocate requests for objects among the sites at which the replicas are stored. Often, these two problems are related in that a placement strategy will have important implications for the request allocation strategy, and vice versa." As such, Rabinovich teaches that selection of strategies for storage and object requests are by no means arbitrary and in fact must be carefully considered.

Rabinovich also teaches (col. 1, line 63 through col. 4, line 40) that many of the techniques developed to date are simply inadequate for the purposes to which the teachings of Rabinovich are intended. Thus, Rabinovich teaches away from arbitrary modification of the teachings of Rabinovich by attempting to modify the teachings of Rabinovich, even within the narrow confines of the arena of deployment contemplated by Rabinovich.

As such, Rabinovich teaches away from the proposed combination of references. It is improper to combine a reference when the reference teaches away from the proposed combination, as is explained below in more detail with reference to MPEP §2145(X)(D)(2), entitled "References Cannot Be Combined Where Reference Teaches Away from Their Combination". This MPEP subsection states that: "It is improper to combine references where the references teach away from their combination. *In re Grasselli*, 713 F.2d 731, 743, 218 USPQ 769, 779 (Fed. Cir. 1983)".

Furthermore, the teachings of either Rabinovich or Zakai are rendered unsuitable for their intended purposes by attempting to combine their disclosures.

It is improper to attempt to combine teachings from references if the proposed combination (i) changes the principle of operation of the reference or (ii) renders the teachings of the reference unsuitable for their intended purpose. These legal concepts are explained below in more detail with reference to MPEP §2143.01.

The former legal concept ((i), supra) is summarized with reference to a subsection entitled "THE PROPOSED MODIFICATION CANNOT CHANGE THE PRINCIPLE OF OPERATION OF A REFERENCE". This subsection states that: "If the proposed modification or combination of the prior art would change the principle of operation of the prior art invention being modified, then the teachings of the references are not sufficient to render the claims prima facie obvious. *In re Ratti*, 270 F.2d 810, 123 USPQ 349 (CCPA 1959)."

The latter legal concept ((ii), supra) is summarized with reference to a subsection entitled "THE PROPOSED MODIFICATION CANNOT RENDER THE PRIOR ART UNSATISFACTORY FOR ITS INTENDED PURPOSE" (see Traverse I).

As such, there can be no motivation to combine their teachings, and the result of such combination fails to provide the elements of Applicant's claims. For at least these reasons, the rejection of claims 20 and 38 is improper and should be withdrawn, and claims 20 and 38 should be allowed.

Traverse V.

Claims 21 and 34 stand rejected as being unpatentable over Zakai in view of Mukherjee. Claim 21 recites "One or more computer readable media as recited in claim 19, wherein the one or more processors are part of a device in a serverless distributed file system", while claim 34 recites "A serverless distributed file system as recited in claim 33, wherein the second plurality of computing devices comprise one or more directory groups", which recitations are not taught, disclosed, suggested or motivated by Zakai, with or without Mukherjee.

Zakai describes swapping of entire physical volumes between storage devices. Mukherjee describes data storage concerns relative to mixed media files. The passing mention of serverless file storage systems in Mukherjee fails to cure the deficiencies of Zakai. Zakai and Mukherjee are concerned with different and distinct issues, and there is no hint whatsoever that might lead one of skill in the art to attempt to combine or modify their teachings as suggested in the Office Action.

As such, there can be no motivation to combine their teachings, and the result of such combination fails to provide the elements of Applicant's claims. For at least these reasons, the rejection of claims 21 and 34 is improper and should be withdrawn, and claims 21 and 34 should be allowed.

Traverse VI.

Claims 42-44 and 46-48 stand rejected as being unpatentable over Falls in view of Zakai. Falls is directed (Title) to: "Transaction synchronization in a disconnectable computer and network". Falls describes: "A method and apparatus are disclosed for synchronizing transactions in a disconnectable network. Each transaction includes operations that were performed on a database replica on one computer while that computer was disconnected from another computer and hence from that other computer's replica. Transaction synchronization, which occurs after the computers are reconnected, transfers information from each computer to the other computer and applies updates to both replicas as appropriate. Transaction logs and clash handling tools may be used with the invention." (Abstract).

In contrast, claim 42 recites "A method, implemented in a directory group, the method comprising: selecting another directory group to participate with in a replica placement improvement process; selecting a file maintained by the directory group; determining whether exchanging a replica of the file with a replica of another file maintained by the other directory group will increase a combined file availability of the files; and having the replica of the file and the replica of the other file exchanged if exchanging the replicas will increase the combined file availability of the files", which is not taught, disclosed, suggested or motivated by the cited references.

The Office Action states (page 18, item 8) that: "As to claim 42, Falls teaches the claimed limitations: "selecting another directory group to participate with in a replica placement improvement process" as [sic] (col. 3, lines 65-67; col. 4, lines 1-5); "selecting a file maintained by the directory group" as [sic] (col. 3,

lines 65-67; col. 4, lines 1-5)". Applicant disagrees and respectfully requests reconsideration in view of the remarks to follow.

The cited portion of Falls states that: "Selected portions of the database may be copied from the network computer to the mobile computer's storage device prior to disconnection as a basis for the mobile computer's virtual network. Copying is accomplished using a device controller in each computer, a replica manager on each computer, and the network link. The device controller on each computer communicates with that computer's storage device to control data transfers."

Falls defines a database at col. 7, line 24 et seq., stating that:

> A database is a collection of related objects. Each object has associated attributes, and each attribute assumes one or more values at any given time. Special values are used internally to represent NULL, NIL, EMPTY, UNKNOWN, and similar values. Each object is identified by at least one "key." Some keys are "global" in that they are normally unique within the entire database; other keys are "local" and are unique only within a proper subset of the database. A database is "hierarchical" if the objects are related by their relative position in a hierarchy, such as a file system hierarchy. Hierarchies are often represented by tree structures.
>
> The target database includes file descriptor objects, directory descriptor objects, directory services objects, printer job objects, or other objects. The target database is distributed in that entries are kept in the replicas 56 on different computers 40. Each replica 56 in the target database contains at least some of the same variables or records as the other replicas 56. The values stored in different replicas 56 for a given attribute are called "corresponding values." In general, corresponding values will be equal.

To equate the database of Falls to the directory group of claim 42 gives the term "directory" a meaning repugnant to the ordinary meaning of the term. The database taught by Falls clearly includes elements inapposite to a directory group, such as printer job objects. Falls also clearly includes directory descriptor objects

as elements reasonably found in a database. It is inappropriate to employ a term in a manner repugnant to the ordinary meaning of the term, as is explained below in more detail with reference to MPEP §2111.01, entitled "Plain Meaning".

This MPEP section states that "THE WORDS OF A CLAIM MUST BE GIVEN THEIR "PLAIN MEANING" UNLESS THEY ARE DEFINED IN THE SPECIFICATION". This MPEP section further states that: "This means that the words of the claim must be given their plain meaning unless applicant has provided a clear definition in the specification. *In re Zletz*, 893 F.2d 319, 321, 13 USPQ2d 1320, 1322 (Fed. Cir. 1989)". The terms "database" and "directory" have different, well-established, plain meanings, as the terms are employed in the vernacular of the relevant arts.

Zakai is void of the term "database" and uses the term "directory" exactly once, stating (col. 3, line 25 et seq.) that: "The cache memory includes a cache memory manager for managing cache accesses and a cache index directory for identifying data stored in the cache." As such, it is inconceivable that combining the teachings of Falls and Zakai could provide the elements recited in claim 42.

Further, Falls teaches away from the teachings of Zakai. More specifically, Falls teaches (col. 2, line 25 et seq.) that:

> However, some of these approaches to replication are not transactional. A transaction is a sequence of one or more operations which are applied to a replica on an all-or-nothing basis. Non-transactional approaches may allow partially completed update operations to create inconsistent internal states in network nodes. Non-transactional approaches may also require a synchronization time period that depends directly on the total number of files, directories, or other objects in the replica. This seriously degrades the performance of such approaches when the network connection used for synchronization is relatively slow, as many modem or WAN links are.

> Moreover, in some conventional approaches potentially conflicting changes to a given set of data are handled by simply applying the most recent change and discarding the others. **Another drawback of several conventional approaches to replication is the requirement they impose that either or both computer systems be locked out of use while the replicas are being synchronized**.

Zakai teaches a system whereby physical volumes are swapped between storage devices, rather than updating of database entries, as in Falls. As such, the physical volumes are not accessible for other purposes during the interval while they are being swapped.

Accordingly, there can be no motivation to combine their teachings, and the result of such combination fails to provide the elements of Applicant's claims. For at least these reasons, the rejection of claims 42-44 and 46-48 is improper and should be withdrawn, and claims 42-44 and 46-48 should be allowed.

Traverse VII.

Claim 45 stands rejected as being unpatentable over Falls in view of Zakai and further in view of Rabinovich. In contrast to the cited references, claim 45 recites "A method as recited in claim 42, wherein having the replica of the file and the replica of the other file exchanged comprises communicating with a first device on which the replica of the file is stored and a second device on which the replica of the other file is stored to have the first device transfer the replica of the file to the second device and delete the replica of the file on the first device, and to have the second device transfer the replica of the other file to the first device and delete the replica of the other file on the second device", which is not taught, disclosed, suggested or motivated by the cited references.

As noted above (Traverse VI), Falls and Zakai fail to provide the subject matter recited in the base claim. The Office Action states (page 21, item 9) that Zakai further teaches (the underlined portion above) and indicates that such appears at col. 6, lines 45-55. Applicant disagrees. The cited portion is included in col. 6, line 43 et seq., reproduced below:

> At step 58, the service processor 28 pairs off the most and least busy of the storage devices 15-17. Thus, the service processor 28 pairs storage devices 15 and 16 and determines that the storage device 17 will not participate in data swaps.
>
> At step 59, the service processor 28 uses a search methodology to select data swaps that decrease workload imbalances by more than a threshold amount. Since the imbalance between the storage devices 15 and 16 is ten accesses per time slice, only data swaps that reduce the imbalance by at least 2.5 accesses per time slice can be selected if the threshold is 25%. Swapping the data of storage volume A with storage volume D or E will reduce the imbalance between storage devices 15 and 16 by 1.33 and 2 accesses per time slice, respectively. Both reductions are too small for the service processor 28 to select these data swaps. Swapping storage volumes A and E will reduce the imbalance between the storage

> devices 15 and 16 by 4.66 accesses per time slice, which is greater than the exemplary threshold of twenty-five percent. Thus, the service processor 28 will select the data swap between volume A and volume E at step 59.
>
> After selecting the data swap between storage volumes A and E, the service processor 28 searches for other data swaps between the remaining storage volumes B, C and D, F. Any further selections of data swaps must further decrease the workload imbalance between the storage devices 15 and 16 by an above threshold amount. Swapping the storage volumes A and E makes the workload of the storage device 15 less than the workload of the storage device 16. Any other data swaps between the storage volumes B, C and the storage volumes D, F will increase the imbalance between the storage devices 15 and 16. Thus, the method 56 will only select to swap the storage volumes A and E for the workloads shown in FIG. 4A.

Zakai describes swapping of entire physical storage volumes (e.g., one of A, B etc. for one of D, E etc., Fig. 1A) in order to balance accessions of distinct memory devices (e.g., 15, 16, 17). This is clearly set forth at col. 5, line 60 et seq., stating that:

> For each identified pair of storage devices 15-17, the service processor 28 searches for swaps of physical storage volumes A-G that produce better balanced workloads (step 59). A swap must reduce historical imbalances of workloads for the paired storage devices 15-17 by more than predetermined threshold amount, e.g., ten percent, to be retained. **Two storage volumes A-G must have the same size and emulation characteristics to qualify as potentially swappable**.

As such, there can be no motivation to combine their teachings, and the result of such combination fails to provide the elements of Applicant's claims. For at least these reasons, the rejection of claim 45 is improper and should be withdrawn, and claim 45 should be allowed.

Unpatentability:

All of the unpatentability rejections (discussed hereinabove with respect to Traverses I through VII) fail to meet the standards set forth in the MPEP for establishing a prima facie case of unpatentability. These are set forth in MPEP §2142, entitled "Legal Concept of Prima Facie Obviousness" (see also MPEP §706.02(j).).

This MPEP section states that "To establish a prima facie case of obviousness, three basic criteria must be met. First, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings." The references fail to teach or disclose the elements recited in the claims. Accordingly, the references cannot provide motivation to modify their teachings to arrive at the invention as claimed, and the Examiner has identified no such teaching or disclosure in the references. As a result, the first prong of the test cannot be met.

MPEP §2142 further states that "Second, there must be a reasonable expectation of success. Finally, the prior art reference (or references when combined) must teach or suggest all the claim limitations."

Inasmuch as the references fail to provide all of the features recited in Applicant's claims, as set forth with specificity in the above traverses, the third prong of the test is not met. As a result, there cannot be a reasonable expectation of success. As such, the second prong of the test cannot be met.

MPEP §2142 additionally states that "The teaching or suggestion to make the claimed combination and the reasonable expectation of success must both be

found in the prior art, not in applicant's disclosure. *In re Vaeck*, 947 F.2d 488, 20 USPQ2d 1438 (Fed. Cir. 1991)." This fourth criterion cannot be met because the references fail to teach or disclose the elements recited in the claims.

Accordingly, the unpatentability rejections fail all of the criteria for establishing a prima facie case of obviousness as set forth in the MPEP.

Inasmuch as there is no guidance within the references, and as there is no basis for the Examiner's contentions within the cited references, the only possible motivation for these contentions is hindsight reconstruction via an improper "obvious to try" standard for finding unpatentability, wherein the Examiner is utilizing Applicant's own disclosure to construct a reason for combining the cited references. The Examiner is reminded that hindsight reconstruction is not an appropriate basis for a §103 rejection. (*See, e.g., Interconnect Planning Corp. v. Feil*, 227 USPQ 543, 551 (Fed. Cir. 1985); *In re Mills*, 16 USPQ2d 1430 (Fed. Cir. 1990) (explaining that hindsight reconstruction is an improper basis for rejection of a claim).).

The impropriety of "obvious to try" as a standard for unpatentability is described in more detail below with reference to MPEP §2145(X)(B). This MPEP section states that:

> The admonition that 'obvious to try' is not the standard under §103 has been directed mainly at two kinds of error. In some cases, what would have been 'obvious to try' would have been to vary all parameters or try each of numerous possible choices until one possibly arrived at a successful result, where the prior art gave either no indication of which parameters were critical or no direction as to which of many possible choices is likely to be successful.... In others, what was 'obvious to try' was to explore a new technology or general approach that seemed to be a promising field of experimentation, where the prior art gave only general guidance as to the particular form of the claimed invention or how to achieve it.

In re O'Farrell, 853 F.2d 894, 903, 7 USPQ2d 1673, 1681 (Fed. Cir. 1988) (citations omitted).

No indication as to which parameters are critical and no direction as to which of many possible choices is likely to be successful has been identified in the references relied upon.

Further, Applicant notes that no evidence has been provided as to why it would be obvious to combine or modify the teachings of these references. Evidence of a suggestion to combine or modify may flow from the prior art references themselves, from the knowledge of one skilled in the art, or from the nature of the problem to be solved. However, this range of sources does not diminish the requirement for actual evidence. Further, the showing must be clear and particular. See *In re Dembiczak*, 175 F.3d 994, 998 (Fed. Cir. 1999).

Conclusion

To recapitulate the several legal arguments relative to alleged unpatentability:

(i) the references teach away from one another and the claimed subject matter; (ii) the teachings of the references are rendered unsuitable for their intended purposes if modified to arrive at the claimed subject matter; (iii) the proposed combinations fail to provide the claimed subject matter; (iv) the rejections fail to meet the criteria for a prima facie showing of unpatentability set forth in the MPEP; (v) the rejections are based on impermissible hindsight; (vi) the rejections employ an improper "obvious to try" standard; and (vii) no proper evidence of suggestion to modify or combine has been provided.

Claims 1-48 are in condition for allowance. Applicant respectfully requests reconsideration and issuance of the subject application. Should any matter in this case remain unresolved, the undersigned attorney respectfully requests a telephone conference with the Examiner to resolve any such outstanding matter.

Respectfully Submitted,

Date: 1/13/05

By: _____
Daniel L. Hayes
Reg. No. 34,618
(509) 324-9256